


# Flexible I/O for the Dynamic Data Center

Mellanox 10/25/50/56/100 Gigabit Ethernet Converged Network Adapters



Mellanox continues its leadership  
in providing high-performance networking technologies  
by delivering superior productivity and scaling in  
10/25/50/56/100 Gigabit Ethernet Adapters,  
enabling data centers to do more with less.

## VALUE PROPOSITIONS

- Data centers deploying virtualized servers see higher demands on the I/O infrastructure. With support for 10/25/50/56/100 Gb/s and up to 512 virtual functions, Mellanox adapters can satisfy the bandwidth demands of a virtualized environment. And with backwards compatibility, the Ethernet cards can be deployed in early generation fabrics giving IT managers an upgrade migration path to 40, 50 and 100 GbE.
- Web 2.0 and cloud service and infrastructure providers need high bandwidth virtualization support and CPU offloads to achieve the highest productivity from their data centers. Mellanox Ethernet adapters support virtualization, overlay networks, CPU offloads and RDMA over Converged Ethernet (RoCE) enabling data center efficiency and scalability.
- Financial institutions utilizing high frequency trading or data exchange applications require low latency fabrics. Mellanox Ethernet adapters deliver low latency Sockets and RDMA solutions for the ultimate in performance optimization.

## **Mellanox 10/25/40/56/100 Gigabit Ethernet Network Interface Cards (NIC)**

deliver high bandwidth and industry-leading connectivity for performance-driven server and storage applications in the most demanding data centers, public and private clouds, Web2.0 and Big Data applications, as well as High-Performance Computing (HPC) and Storage systems. Clustered databases, web infrastructure, and high frequency trading are just a few example applications that will achieve significant throughput and latency improvements resulting in faster access, real-time response and increased number of virtual machines per server.

## **World-Class Ethernet Performance**

Mellanox Ethernet adapters utilizing IBTA RoCE and RoCE v2 technology provide efficient RDMA services, delivering high performance to bandwidth and latency sensitive applications. With link-level interoperability within existing Ethernet infrastructure, Network Administrators can leverage existing data center fabric management solutions.

Applications utilizing TCP/UDP/IP transport can achieve industry-leading throughput over 10/25/40/56/100GbE. The hardware-based stateless offload and flow steering engines in Mellanox adapters reduce the CPU overhead of IP packet transport, freeing more processor cycles to work on applications and manage more Virtual Machines, where applicable. Sockets acceleration software further increases performance for latency sensitive applications.

## **I/O Virtualization**

Mellanox Ethernet adapters provide dedicated adapter resources and guaranteed isolation and protection for virtual machines (VM) within the server. Mellanox adapters gives data center managers better server utilization and LAN and SAN unification while reducing cost, power, and cable complexity.

## **Overlay Networks**

New large-scale clouds require the implementation of overlay network protocols in order to overcome the issues of security and isolation within the cloud and the limitations of VLAN. Mellanox Ethernet adapters with hardware offload capability for VXLAN and NVGRE brings a unique value to cloud providers enabling them to reduce their CPU overhead and consequently to reduce their OPEX and CAPEX by supporting more cloud tenants over the same infrastructure . Moreover, ConnectX-4 hardware capability for encapsulating and decapsulating the Overlay Networks' protocol headers further improves the utilization of the Cloud servers.

## Quality of Service

Resource allocation per application or per VM is provided and protected by the advanced QoS supported by Mellanox adapters. Service levels for multiple traffic types can be based on IETF DiffServ or IEEE 802.1p/Q, along with the DCB enhancements, allowing system administrators to prioritize traffic by application, virtual machine, or protocol. This powerful combination of QoS and prioritization provides the ultimate fine-grain control of traffic – ensuring that applications run smoothly in today's complex environment.

## Coherent Accelerator Processor Interface (CAPI)

With its support for CAPI, ConnectX-4 EN provides the best performance for Power and OpenPower based platforms. Such platforms benefit from better interaction between the Power CPU and the ConnectX-4 EN adapter, lower latency, higher efficiency of storage access, and better Return on Investment (ROI), as more applications and more Virtual Machines run on the platform

## Complete End-to-End Ethernet Networking

Mellanox adapters are part of a full 10/25/40/56/100 Gigabit Ethernet end-to-end portfolio for data centers, which includes switches, application acceleration packages, and cables. Mellanox's SwitchX family of Ethernet switches and Unified Fabric Management software incorporate advanced tools that simplify networking management and installation, and provide the needed capabilities for the highest scalability and future growth. Mellanox's messaging and storage acceleration packages deliver additional capabilities for the ultimate server performance. With Mellanox end to end, IT managers can be assured of the highest performance, most efficient network fabric.


## BENEFITS

- *Improved productivity and efficiency*
- *Smart interconnect for x86, Power, ARM, and GPU-based compute and storage platforms Industry-leading throughput and latency performance*
- *Enabling I/O consolidation by supporting TCP/IP, Storage and RDMA over Ethernet transport protocols on a single adapter*
- *Supports industry-standard SR-IO Virtualization technology and delivers VM protection and granular levels of QoS to applications*
- *High-availability and high-performance for data center networking*
- *Cutting-edge performance in virtualized overlay networks (VXLAN and NVGRE)*
- *Increased VM per server ratio*
- *Software compatible with standard TCP/UDP/IP and iSCSI stacks*
- *High level silicon integration and no external memory design provides low power, low cost and high reliability*

## TARGET APPLICATIONS


- *Public and private clouds*
- *Data analytics applications Web 2.0 data centers and cloud computing*
- *Data center virtualization*
- *Low latency financial services*
- *I/O consolidation (single unified wire for networking, storage and clustering)*
- *Video streaming*
- *Enterprise data center applications*
- *Accelerating back-up and restore operations*

## ConnectX<sup>3</sup>




Ports	1 x 1/10GbE	2 x 1/10GbE	1 x 1/10/40/56GbE	2 x 1/10/40/56GbE
Connector	SFP+	SFP+	QSFP	QSFP
Cabling Type*	Direct Attached Copper SR and LR Fiber Optic			
Host Bus	PCIe 3.0			
Features	Stateless Offload, RDMA, SR-IOV, DCB, Precision Time Protocol			
OS Support	RHEL, SLES, Microsoft Windows Server 2008/2012, FreeBSD, Ubuntu 12.04, VMWare ESXi 4.x/5.x			
Ordering Number	MCX311A-XCAT MCX341A-XCAN MCX341A-XCBN	MCX312A-XCBT MCX342A-XCAN	MCX313A-BCBT	MCX314A-BCBT

## ConnectX<sup>3</sup> PRO



Ports	1 x 1/10GbE	2 x 1/10GbE	1 x 1/10/40/56GbE	2 x 1/10/40/56GbE
Connector	SFP+	SFP+	QSFP	QSFP
Cabling Type*	Direct Attached Copper SR and LR Fiber Optic			
Host Bus	PCIe 3.0			
Features	Stateless Offload, VXLAN and NVGRE Offload, RDMA, SR-IOV, DCB, Precision Time Protocol			
OS Support	RHEL, SLES, Microsoft Windows Server 2008/2012, FreeBSD, Ubuntu 12.04, VMWare ESXi 4.x/5.x			
Ordering Number	MCX311A-XCCT	MCX312B-XCCT	MCX313A-BCCT	MCX314A-BCCT

## ConnectX<sup>4</sup>



Ports	1	2	1	2
Speed	40/56GbE	40/56GbE	100GbE	100GbE
Connector	QSFP			
Host Bus	PCI Express 3.0 x16			
Features	RoCE, GPUDirect, SR-IOV, Overlay Networks, Stateless Offloads, Signature Handover, Dynamically Connected Transport			
OS Support	RHEL, CentOS, SLES, OEL, Windows, ESX/vSphere, Ubuntu, Citrix, Fedora, FreeBSD			
Ordering Number	MCX415A-B	MCX416A-B	MCX415A-C	MCX416A-C

*\*Please visit Mellanox's web site for more cable information, best usage practice and availability.*

## FEATURE SUMMARY

### ETHERNET

- IEEE Std 802.3ba 40 Gigabit Ethernet
- IEEE Std 802.3ae 10 Gigabit Ethernet
- IEEE 802.3az Energy Efficient Ethernet
- IEEE 802.3ap based auto-negotiation and KR startup
- Proprietary Ethernet protocols (20/40GBASE-R2, 50/56GBASE-R4)
- IEEE Std 802.3ad, 802.1AX Link Aggregation
- IEEE 802.1Q, .1p VLAN tags and priority
- IEEE Std 802.1Qau (QCN) Congestion Notification
- IEEE P802.1Qaz (ETS) Enhanced Transmission Selection
- IEEE P802.1Qbb (PFC) Priority-based Flow Control
- IEEE 802.1Qbg
- Multicast
- Jumbo frame support (9.6KB)
- 128 MAC/VLAN addresses per port

### TCP/UDP/IP STATELESS OFFLOAD

- TCP/UDP/IP checksum offload
- TCP Large Send (< 64KB) or Giant Send (64KB-16MB) Offload for segmentation

- Receive Side Scaling (RSS) up to 32 queues

- Line rate packet filtering

### ADDITIONAL CPU OFFLOADS

- RDMA over Converged Ethernet (RoCE and RoCE v2)
- Traffic steering across multiple cores
- Intelligent interrupt coalescence
- Stateless offloads for overlay networks and tunneling protocols
- Hardware offload of encapsulation and decapsulation of NVGRE and VXLAN overlay networks

### HARDWARE-BASED I/O VIRTUALIZATION

- Single Root IOV
- Address translation and protection
- Dedicated adapter resources
- Multiple queues per virtual machine
- VMware NetQueue support

### REMOTE BOOT

- Remote boot over Ethernet
- Remote boot over iSCSI
- PXE and UEFI

## COMPLIANCE

### SAFETY

- USA/Canada: cTUVus UL
- EU: IEC60950
- Germany: TUV/GS
- International: CB Scheme

### EMC (EMISSIONS)

- USA: FCC, Class A
- Canada: ICES, Class A
- EU: EN55022, Class A
- EU: EN55024, Class A
- EU: EN61000-3-2, Class A
- EU: EN61000-3-3, Class A
- Japan: VCCI, Class A
- Australia: C-Tick
- Korea: KCC

### ENVIRONMENTAL

- EU: IEC 60068-2-64: Random Vibration
- EU: IEC 60068-2-29: Shocks, Type I / II
- EU: IEC 60068-2-32: Fall Test

### OPERATING CONDITIONS

- Operating temperature: 0 to 55° C
- Air flow: 100LFM @ 55° C
- Requires 3.3V, 12V supplies

## COMPATIBILITY

### PCI EXPRESS INTERFACE

- PCIe Base 3.0 compliant, 1.1, 2.0 compatible
- Fits x8 or x16 slots
- Support for MSI/MSI-X mechanisms
- Coherent Accelerator Processor Interface (CAPI)

### CONNECTIVITY

- Interoperable with 10/25/40/50/100Gb Ethernet switches
- Passive copper cable with ESD protection
- Powered connectors for optical & active cable support
- QSFP to SFP+ connectivity through QSA module

### OPERATING SYSTEMS/DISTRIBUTIONS

- RHEL/CentOS
- Windows
- FreeBSD
- VMware
- OpenFabrics Enterprise Distribution (OFED)
- OpenFabrics Windows Distribution (WinOF)

### MANAGEMENT

- MIB, MIB-II, MIB-II Extensions, RMON, RMON 2
- Configuration and diagnostic tools

### GENERAL

- Adapters for Open Compute Project (OCP)
- Adapters with combined UEFI/Legacy ROM

\* Product images may not include heat sync assembly; actual product may differ.

© Copyright 2014, Mellanox Technologies. All rights reserved.  
Mellanox, BridgeX, ConnectX, CORE-Direct, InfiniBridge, InfiniHost, InfiniScale, MLNX-OS, PhyX, SwitchX, UFM, Virtual Protocol Interconnect and Voltaire are registered trademarks of Mellanox Technologies, Ltd. Connect-IB, CoolBox, FabricIT, Mellanox Federal Systems, Mellanox Software Defined Storage, MetroX, Open Ethernet, ScalableHPC and Unbreakable-Link are trademarks of Mellanox Technologies, Ltd. All other trademarks are property of their respective owners.



350 Oakmead Parkway, Suite 100  
Sunnyvale, CA 94085  
Tel: 408-970-3400  
Fax: 408-970-3403  
www.mellanox.com

3530BR Rev 6.6